Paper Number 211

DESIGN CONSIDERATIONS FOR THE MARS NETWORK OPERATIONS CONCEPT

R. J. Cesarone, S. N. Rosell, S. A. Townes, D. J. Bell, D. S. Abraham, and G. J. Kazz
Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

NASA's Mars Surveyor Program is developing plans for an evolving constellation of satellites in Mars orbit, called the Mars Network, to support the telecommunication and navigation needs of future Mars exploration. As an extension of the Deep Space Network (DSN), this network must support a diverse set of users, including planned missions, as well as yet-to-be-defined mission concepts. Network elements comprise low-cost micro-satellites, as well as Mars Areostationary Relay Satellites (MARSAT). Mars Network operations must provide an efficient, and largely autonomous, switching mechanism to coordinate the demands of the burgeoning number of user assets expected at the planet. In addition, operations of the network assets themselves must also be done in as efficient and autonomous a manner as possible. The operations concept considers the topology of the network and its evolution, due to planned enhancements and unplanned anomalies. Other operational issues, such as data management, demand vs. scheduled access, prioritization and qualities of service, handoffs and multiple satellites in view will be discussed.

INTRODUCTION

During the summer of 1998, a NASA/JPL Mars Exploration Program Architecture Definition Study made a strong recommendation for a low-cost *in-situ* communications and navigation satellite network to provide enabling and enhancing support for the international exploration of Mars. This would comprise the first element of an "Interplanetary Internet" and initiate a "virtual presence throughout the solar system" as called for in NASA's Strategic Plan.[1] FY99 funding was provided to conduct a phase A study of this Mars Network. More recently, following the demise of the Mars Climate Orbiter and Mars Polar Lander, various review boards and teams re-examining the Mars Exploration Program Architecture have voiced their support for such an *in-situ* network. As of the time of this writing, implementation activities have been included in the President's FY01 budget submittal. Although this submittal must work its way through the usual Congressional budgetary process, the outlook is positive. Consequently, FY00 ramp-up funds have been made available for ongoing work at JPL.

Implementation of the network will begin with the gradual deployment of a constellation of relatively low-altitude, low-cost micro-satellites (MicroSats).[2, 3, 4] These spacecraft are based on the use of a common micromission bus designed for launch on the Ariane 5 as a secondary payload.[5, 6] The first network MicroSat will be dispatched to Mars in the 2003 opportunity. One or two additional MicroSats are to be launched on each succeeding Earth-to-Mars opportunity, every 26 months, until the proposed steady-state constellation of six satellites is achieved. The MicroSats will be augmented, as needed, with more capable Mars areostationary relay satellites (MARSATs) to support future robotic outposts and ensuing human missions. The Mars Network will be deployed, maintained and operated as an extension of NASA's Deep Space Network (DSN) to serve as reliable communications and navigation infrastructure available for the support of future Mars missions.

This paper discusses specific design considerations that influence the Operations Concept for the Mars Network. Operational functions of individual network elements are covered first, followed by a discussion of network-level functions. The operations concept addresses the topology of the network spanning the diverse number of switched links among the Mars surface, Mars orbit, Mars approach trajectory, Earth and inter-network cross-links. Evolution of this topology, due to planned enhancements and unplanned anomalies, is discussed. Other operational issues, peculiar to network-level functionality, will be discussed. Finally, an overview of applicable protocols is provided.

NETWORK ELEMENTS

The Mars communications and navigation infrastructure, depicted in Figure 1, comprises four main elements. The first of these is a set of Mars-orbiting, relatively low-altitude micro-satellites (MicroSats). The currently envisioned MicroSats are to be launched as piggyback payloads on the Ariane 5 launch vehicle. Despite the modest size of these assets, they are able to provide noticeable improvements in connectivity and end-to-end data rates, as well as the ability to enable position determination in a manner analogous to that of the Global Positioning System (GPS) at Earth.

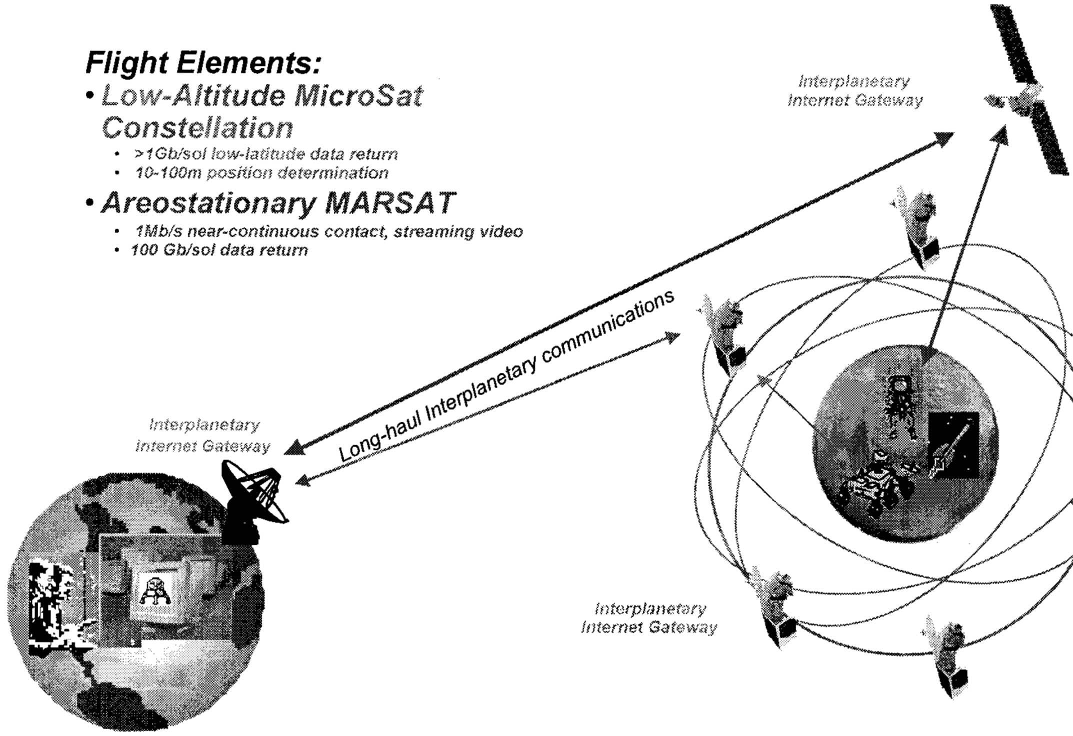

Figure 1 Mars Network Overview

A first MicroSat is to depart for Mars in the 2003 opportunity, to eventually take up residence in an 800-km, near-equatorial orbit. At each succeeding Earth-to-Mars opportunity (~ 26 months), two more such spacecraft are proposed to be dispatched to Mars, targeted for near-equatorial and high-inclination orbits as needed. Equatorial orbiters provide excellent connectivity to low-latitude landed-elements, which are expected to include most sample return operations. Highly inclined orbiters round out the constellation by providing global coverage for the benefit of high-latitude surface elements. Six satellites are nominally planned for the steady-state constellation.

The second element consists of a small number of Mars-orbiting areostationary satellites (MARSATs). ΔV requirements to attain this high-altitude orbit necessitate that these be heavier, more expensive, prime launch vehicle payloads and hence limited in number. Nevertheless, they will provide dramatic increases in end-to-end data rates, with nearly continuous coverage over most of the Martian hemisphere under their stationary longitudes. The first of these assets will hopefully launch at the 2005 Earth-to-Mars opportunity, at the earliest. They will provide the high-capacity link that will be required as the near-term robotics program transitions to robotic outposts and then to the set of missions culminating in humans on Mars. The necessary equatorial orbit and static link geometry will lessen the utility of the areostationary satellites for global positioning.

A third element of the architecture is the set of large deep space tracking antennas located on Earth. These will primarily comprise the antennas of the DSN, located in California, Spain and Australia. However, tracking assets of other nations are expected to interoperate with the DSN to expand capacity and to support the overall effort.

The fourth element is the set of systems and software that tie the whole architecture together, and provide a front-end through which the public can participate in the Martian adventure. Indeed, the total system can be thought of as an extension of DSN nodes and services to the Mars *in-situ* region. The concept has been likened to the beginnings of an interplanetary Internet that will bring the exploration of Mars right into our living rooms.

USER MISSION SET

A number of Mars exploration missions are candidates to receive Mars Network support over the next decade. While the specific mission architecture for Mars exploration is still evolving, most of the architecture options include the following exploration elements:

- Orbiters – Perform one or more of the following functions: delivery of landers and/or probes to the Martian surface, orbital science and site reconnaissance, collection and return of sample canisters to Earth, and communication relay. Most orbiters will be equipped with both X-band direct-to-Earth and UHF orbit-to-Mars surface communications capability.

- Landers – Serve both as platforms for conducting *in-situ* science and technology experiments and as staging areas for the collection and Earth return of samples. Among the former are two European elements: Beagle 2, which is delivered to Mars during the 2003 opportunity for a 180-day mission to search for life at a site within 0 to 10 deg latitude, and Netlanders, four stations, doing seismic, climate and other global investigations, delivered during the 2005 opportunity for one year of surface operation at sites within ± 35 deg latitude. Between 2003 and 2009, sample return lander missions will occur in Mars' equatorial zone with one or more employing a Mars Ascent Vehicle (MAV) to launch samples into Mars orbit for return to Earth. To facilitate such activities, Doppler surface location determination for these missions is desired to within 1 km. Landers will be equipped with both X-band direct-to-Earth and UHF Mars surface-to-orbit communications capability.

- Rovers – Enable surface exploration and sample acquisition well beyond the confines of the landers that deliver them. Rover missions are under study which involve traverses of up to several kilometers from the delivery lander. To facilitate such traverses, Doppler surface location for these missions is desired to within significantly less than 1 km. Rovers will be equipped with a two-way UHF radio to communicate with the delivery lander and/or orbiters.

- Canisters – Contain the samples gathered by the landers and rovers and ride on the lander-launched MAVs. The MAVs inject the canisters into a 600-km altitude, 45-deg inclined parking orbit for later retrieval by an Orbiter/Earth Return Vehicle. Each canister will be equipped with a low power UHF transponder, which will provide a continuous Doppler signal that can be received by the Mars Surveyor '01 Orbiter, the European Space Agency (ESA)/ Agenzia Spaziale Italiana (ASI) Mars Express Orbiter, and a MicroSat for orbit determination.

- Scouts – Precede larger landers and provide reconnaissance data on potential landing sites. These small probes will utilize UHF links with orbiting assets to provide data during descent and, hopefully, for up to 7 sols of surface operations. As with the landers and rovers, fairly precise Doppler location determination is essential.

- Micromission Probes – Look like MicroSats, but carry science payloads rather than telecom payloads. These payloads can be Deep Space 2-like probes, aerobots, or even airplanes, assuming they can fit within the Micromission spacecraft's small mass and volume payload constraints. All such payloads will carry either one-way or two-way UHF surface-to-orbit communications capability; and, some may require Doppler location determination.

One possible scenario for the deployment of these exploration elements appears in Figure 2.

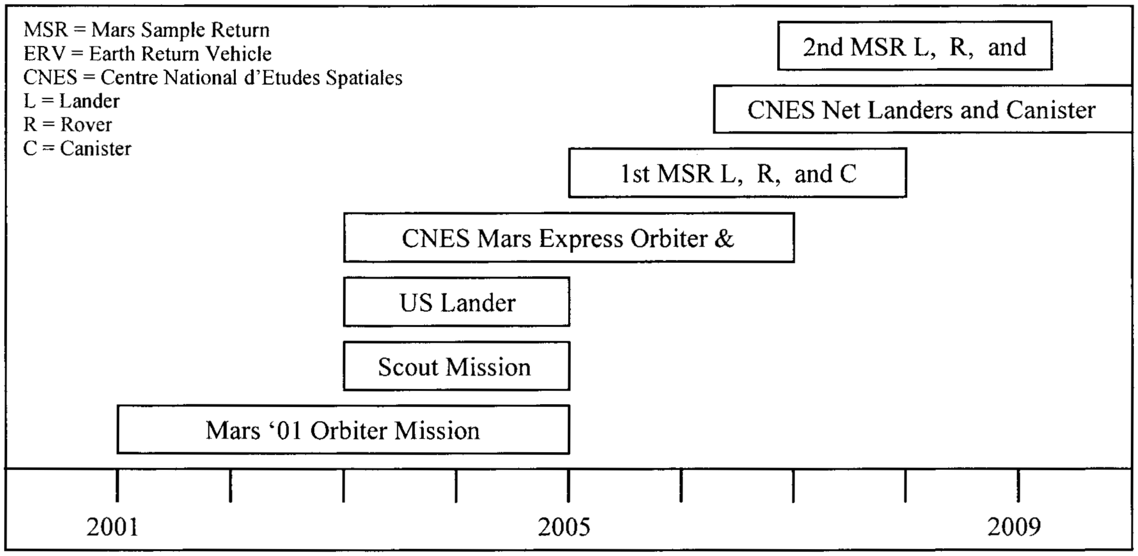

Figure 2 Possible Scenario for the Mars Exploration Missions

ELEMENT LEVEL FUNCTIONS

Before discussing network functionality, it is worth noting that individual satellites of the Mars Network constellation can independently provide three key functions or services. The first, and most obvious, is point-to-point communications. This entails the transmission of both command data on the forward link, generally from Earth to the user asset, as well as telemetry data on the return link in the opposite direction. The second service provides radiometric data usable for surface positioning or flight navigation. Data types will be based on 2-Way Doppler and range measurements, though innovative variants of these may also be available. Note that the varying link geometry of the low-altitude MicroSats will make them more useful for navigation purposes than the areostationary MARSATs. Finally, each Mars Network spacecraft will carry an Ultra-Stable Oscillator (USO) accurate to the 5×10^{-13} s/s level. Hence, they will be able to provide time signals usable by user spacecraft at the planet. These services are enabled by the spacecrafts' *in-situ* radio payloads, comprising hardware and software, and also known as Mars Network Nodes.[7] Eventually, these nodes will transition from providing element-level to Network-level services.

EVOLUTION OF MARS NETWORK TOPOLOGY

Figure 3 provides a possible deployment schedule for a network at Mars. In addition to the MicroSats and MARSATs of the actual Mars Network, the figure shows timelines for remote-sensing missions that also carry *in-situ* communications equipment. The constellation of small satellites is shown as extending beyond 2012. Additional areostationary satellites, for support of high data rates and potential manned missions, are depicted as overlapping the small satellites, starting in 2007, though this may occur as early as 2005. Because the network will be emplaced over a series of Earth-to-Mars opportunities, its time to completion will be necessarily lengthy. During this deployment period, the network's topological state will evolve, with connectivity and functionality increasing as additional Network nodes are deployed. However, during each phase of the deployment, Mars Network nodes will provide usable levels of communication, radiometric positioning, and time service. It is instructive to examine some candidate topologies for the fully-implemented architecture. A number of these are shown in Figure 4. Characteristics of each candidate are provided in Table 1.

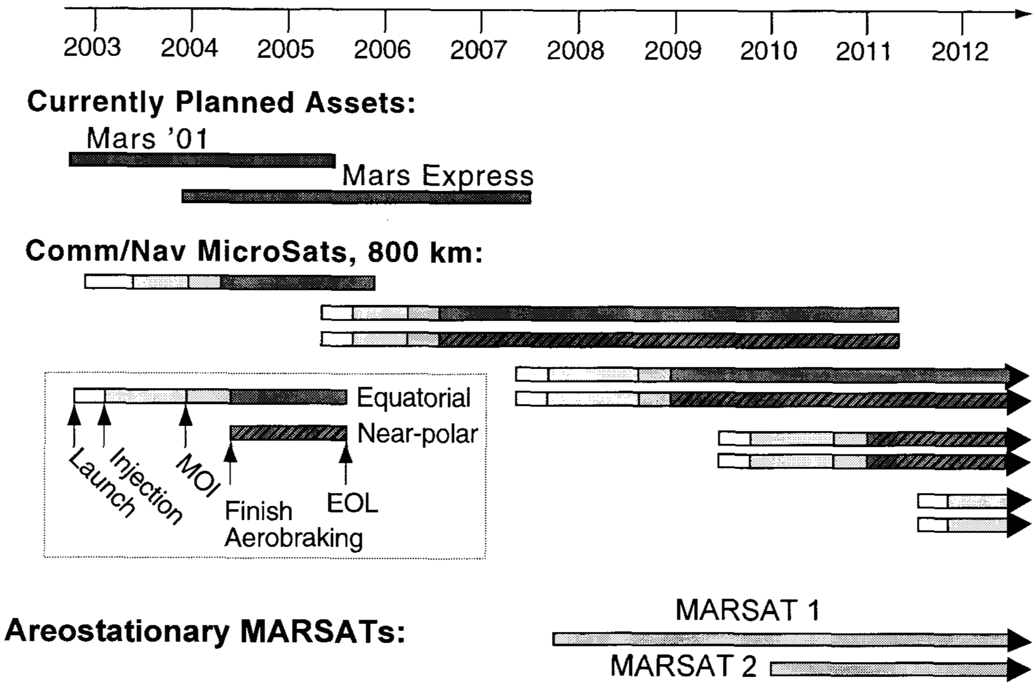

Figure 3 Possible Mars Network Deployment Plan (To Be Updated)

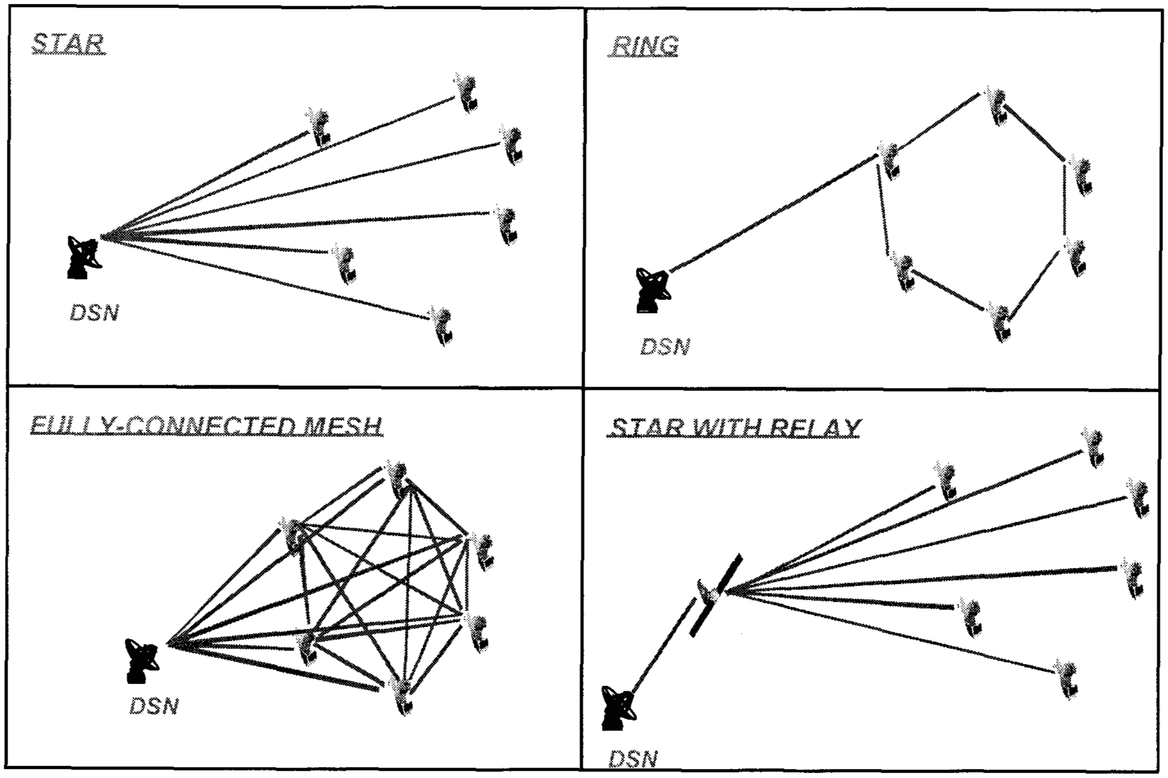

Figure 4 Candidate Mars Network Topologies

Table 1
Characteristics of Candidate Network Topologies

Network Topology	Star	Ring	Fully-Connected Mesh	Modified Star
Flexibility	Limited	Moderate	Very High	Limited
Roles of Network Elements	Similar	Dissimilar	Similar or Dissimilar	Dissimilar
Autonomy	Little	Moderate	High	Moderate
Latency for Network Control	Long	Short	Short	Short
Susceptibility to Outages	Low	High	Low	Low
Occultation Opportunities	Not Available	Available	Available	Very Limited
Management & Operations Complexity	Simple	Moderate	Very High	Low

The first candidate is the classic star configuration. It is the simplest form that allows network-like functionality. The center of the star is NASA's Deep Space Network (DSN), located on Earth. All communications lines route through the DSN. This configuration would probably not be considered very desirable as an end-state topology because of the long light-time delays on each arm which would make inter-satellite communications inefficient. Note also that the links to each satellite are active or not depending upon whether visible from Earth or not. However, its very simplicity may make it an ideal candidate for the initial stages of Mars Network deployment.

The second candidate, based upon another classic configuration, has the satellites at Mars configured in a ring. This arrangement adds the new feature of inter-satellite cross-links among the network nodes. At least one satellite in the ring must act as a gateway to the DSN and probably two or more since the gateway can become hidden as it rotates around Mars. Note also that for some orbit choices, these inter-satellite links may not be active all of the time due to either long ranges or occultation by the planet. Perhaps the key deficiency of the ring configuration is its susceptibility to outages. This can occur when an anomaly affects any node in the network, thus breaking the ring structure, and it is particularly severe when it happens to a gateway node.

The third candidate is called a fully-connected mesh. This is a form that enables a significant increase in operational flexibility. It has all the desirable features expected of a fully operational network, including high autonomy, short latency and low susceptibility to outages. However, all these desirable features come at a price, namely the substantial increase in complexity of network equipment, management and operations.

The fourth candidate is called a star with relay. The intent of this topology is to regain some of the operational simplicity of the prototypical star configuration, but without its inherent limitations. This is achieved by replacing the Earth-based DSN at the center of the star with the Areostationary MARSAT. This enables a great reduction in latency within the network at Mars and a consequent increase in autonomy. Cross-links do not occur between or among MicroSats, though they are available between each MicroSat and the MARSAT. The one weak link in this, or course, is the MARSAT relay failure. This can be overcome with a redundant MARSAT or accepting lower performance on the direct-to-Earth links from the MicroSats.

Phase 1 of the Mars Network will be deployed in the 2003 opportunity. During this phase of the activity there will be little actual network-level functionality. This is due to the fact that there will be only one, or possibly two, MicroSats deployed, and very likely no MARSAT. Thus the available satellite will operate as a relay node in a point-to-point communication system. However, the various techniques and operational approaches that will characterize eventual network operations can be tested and validated during this phase. It is worth noting that two additional orbiters, namely Mars Surveyor '01 and Mars Express, each carrying relay communications capabilities, will also be operational at this time. Though ongoing efforts are striving to assure compatibility among these assets and Mars Network MicroSats, the degree to which network-like interoperability will be achieved is unclear at this time.

Phase 2 will comprise the 2005 and 2007 opportunities and will witness the build-up of an interim satellite constellation. During this phase, four additional MicroSats may be deployed, along with one or possibly two MARSATs. The original MicroSat from the 2003 opportunity will likely reach its end-of-life during this time period. As Phase 2 proceeds, the Mars Network will be able to begin providing actual network-level services.

Phase 3 adds the 2009 opportunity and results in the full constellation of six MicroSats plus two MARSATs. At this point the ability to provide the full suite of network-level services will have been emplaced.

NETWORK-LEVEL FUNCTIONS

As the Mars Network grows and evolves, it will move out of the realm of providing only point-to-point services and move into the realm of full network-like functionality. Prior to this time, it is natural to think of the individual orbital elements as deep space satellites with all the usual systems and subsystems. However, as the Network era dawns, each orbiting asset can be more simply described as a store-and-forward switch in a larger architectural context, in essence a "Big Switch in the Sky."

Switching

With this concept in mind, it is instructive to examine the various elements that these assets must connect by means of this switching function. These include, but are not necessarily limited to, the following links.

1) <u>Earth (DSN) ↔ Mars Landed Element</u>: This is the most obvious connection, that of connecting an end user on Earth with an asset on Mars.

2) <u>Earth (DSN) ↔ Mars Network Satellite</u>: For the sake of its own operations, the Network elements must communicate with controllers on Earth.

3) <u>Mars Landed Element ↔ Mars Landed Element</u>: This type of connection is not likely in the early stages of robotic exploration. However, as human exploration of Mars commences, astronauts will likely require *in-situ* links with deployed assets operating over the horizon. These links may very well have real-time joystick requirements.

4) <u>Mars Network Satellite ↔ Mars Network Satellite</u>: Cross links between or among satellites are useful for various operations of the network itself, e.g., data transfer, rerouting to handle outages, dissemination of orbit ephemeris information, and local time synchronization. In addition, cross-links have the capability of carrying out very useful radio occultation observations. As the line-of-sight connecting two satellites in communications begins to intersect the planet's atmosphere, RF signal attenuation will occur. Monitoring signals during these events can yield valuable insights into the time-varying nature of the Martian atmosphere.

5) <u>Mars Network Satellite ↔ Mars Orbiter</u>: Most user assets in Mars orbit will have their own high data rate links back to Earth via the DSN. However, certain assets may have extremely low power and thus be limited only to localized links. In particular, Mars Sample Return missions are expected to launch a canister from the Martian surface carrying the precious sample cargo. Because of mass constraints, these will be very small and hence have very low power. A further complexity is that they may be stored in Mars orbit for up to 2 years prior to retrieval by an Earth return vehicle. During this period of time, Network satellites will communicate with, and maintain up-to-date ephemeris information for these orbiting canisters.

6) <u>Mars Network Satellite ↔ Approaching/Departing Spacecraft</u>: Because the orbital ephemerides of Network satellites will be very well known, they can be used essentially as navigation beacons for other spacecraft approaching or departing the Mars vicinity. This will enable very precise navigation with obvious benefits for precision landing, mission safety, etc.

Network Access

A single satellite in a point-to-point communications system is very simple to schedule. However, as the number of user and Network assets at Mars both continue to increase, the issue of scheduling becomes more complex. This can be done either by demand or scheduled access. Scheduled access would involve maintaining an up-to-date database of the communications requirements of all user assets and the communications opportunities of all Network assets. With valid information available, scheduling software could then be run at a Network control point and uploaded to Network assets, which would then execute the communications sessions at the appropriate times. Note that the Network control point could be, but is not necessarily, on Earth. For example, scheduling could be done at a MARSAT, which, in addition to its own communications functions, is also acting in this capacity. Demand access, on the other hand, requires no such centralized planning. However, it does require a previously agreed upon set of standards and protocols that determine the establishment or disconnection of links in real time.

It is likely that the early stages of the Mars Network will rely on scheduled access, with that function being performed by an operations team on Earth. As experience, and confidence, in Network operations is gained, there will be an increasing rationale for migrating this scheduling function to the assets themselves. This is possible because network nodes will know their own locations, as well as those of the exploration assets they support. That fact, plus the increasing autonomy capabilities of the Network, in its evolved stages, will enable on-board scheduling to occur. Going one step further, the scheduling function can be replaced by an autonomously operated demand access scheme. One method under consideration for implementing this scheme is by use of a hailing frequency broadcast by a network node. Upon receipt of the hail, exploration assets can request a link. If and when the Network node accepts the request, it then assigns a frequency for the communications session, after which the session is executed.

One of the key issues for autonomous scheduling is prioritization. Although Mars Network satellites will initially be capable of dual, and eventually, several simultaneous links, this capability will have its limits. In the event that there is contention among user assets for a network "circuit" it will be necessary to determine which asset has priority. Schemes that always assign the link to the first requestor may prove too simplistic. Consider a case in which a high-priority user experiencing a critical event has a view period to a network satellite that begins mere seconds after a much lower priority user has grabbed the link. Clearly, situations like this will have to be managed by judicious application of appropriate protocols.

Two other related issues are involved in the scheduling aspect of Network operations. The first is the subject of handoffs. It may occur that a communications session between a surface asset and a Network satellite is insufficient to transfer all the buffered data. In that case a handoff to another Network asset will be needed. Conversely, the surface asset may have more than one Network asset in its field of view. In this case, it will be necessary to make a determination of which satellite should be utilized for the communications session. Handoffs and choice of multiple Network satellites will also be enabled by the use of appropriate protocols.

Quality of Service (QOS)

The Network will ultimately have to transport many different types of data, each perhaps requiring its own appropriate QOS. Some, such as key science data and critical command sequences, will require guaranteed, error-free delivery, with full data accountability and high security. These data types will likely require acknowledged delivery of validated file transfers. On the other hand, they will be able to tolerate a relatively long latency. Other data types, such as routine science data and high-throughput voice and video, may be able to utilize links with significantly less reliability and security. Further, it is more important to acquire voice and video in sequence, and with low latency than it is to ensure complete and accurate transmission.

Data Management

Assuming that Network satellites have cross-link capabilities, it will be possible to develop a data-management scheme. One example of this might occur when a satellite with a full buffer, and needing to receive more data from the surface, transfers its data to another satellite in the constellation which it can see but which the surface asset can not see. Another possibility occurs when a satellite, in the process of downlinking to Earth, enters Earth-occultation by Mars. If it can see, and cross-link to, another constellation satellite, it can continue the downlink session it would otherwise have to terminate.

Protocols

Recently, the Consultative Committee on Space Data Systems (CCSDS) has created the *Proximity-1 Space Link Protocol*.[8] This protocol defines both the physical layer (RF and modulation) as well as the data link layer (frame structure and validation, media access, data services, input/output) interfaces for spacecraft operating at short range. These links are characterized by short time delays, moderate (not weak) signal levels, and short, independent sessions. The protocol specifies the frequencies, modulation, synchronization and framing, coding, addressing, link establishment and termination procedures for point-to-point links. The protocol also supports simultaneous communication between one caller to many responders on the forward link. In addition, the protocol enables a method for determining downlink prioritization amongst several users on the return link. The protocol supports all levels of directionality: simplex, half-duplex, and full-duplex and is therefore applicable to the large variety of spacecraft which the Mars Network assets must support.

This protocol contains two Qualities of Service (QOS): the sequence controlled service provides a "go-back-n" methodology to ensure reliable proximity frame transfer at the data link layer. Expedited service will be utilized by spacecraft to support both expedited commanding (by passing sequence controlled service) and bi-directional file transfer. For file transfer, reliability is ensured at the transport layer by means of the *CCSDS File Delivery Protocol* (CFDP) using a selective repeat methodology.

In order to move beyond the era of stream based telemetry and into the new era of reliable, demand driven telemetry, it is envisioned that missions supported by the Mars Network assets will implement CFDP at the transport layer running under Proximity-1 at the data link layer.

Navigation State Determination

The navigation and positioning capabilities of the Mars Network will also evolve with time. In the early stages, individual Network nodes will provide the radiometric data types. In the later stages, the satellites will work cooperatively, perhaps in real-time, to jointly enable determination of a user's state.[9] There are various ways in which this might be accomplished. Analogous to the Earth-based Global Positioning System (GPS), satellites might broadcast information that user assets can rely on to make their own position determinations. Conversely, the Network may have the capability to compute user position and/or velocity from data received at one or more of its satellites, then transmit this information to the user asset, as needed. In either case, the function requires the coordinated use of the satellites as one integrated network.

Maintain Mars Standard Time

Incorporation of USOs on the Network satellites will enable them to keep accurate time, synchronized at appropriate intervals, to time signals provided by the DSN. However, no two clocks are exactly the same. Thus it is likely that one satellite will be selected to be the keeper of "Mars Standard Time." Most likely this selection will be based upon whichever USO empirically demonstrates the greatest stability. Following this selection, other satellites in the Network can have their clocks updated by the master timekeeper, via cross-links. Finally, accurate time signals can then be provided to the user assets when they have communications sessions with the Network satellites.

ROLES OF OPERATIONS CENTERS

Although Network operations is likely to assume ever increasing autonomy as time progresses, it is virtually certain that some amount of operational activity will still need to take place here on Earth. Since Network elements are indeed fully-functional spacecraft, it is assumed that there will be a Spacecraft Operations Center. This center will monitor items such as spacecraft health, consumables status, system trends, etc. All of the spacecraft functions are only there to enable the actual payload functions, i.e., communications, navigation and time determination. Because the heart of the Network resides in the payload, it may be desirable to have a separate Payload Operations Center, which will monitor the performance of the key services provided. Finally, a Network Operations Center provides management, performance monitoring and, as needed, control of the higher-level Network functions.

All three layers of operations, Spacecraft, Payload and Network, are really only there to support provision of services to a set of end users. These users will have their own Mission Operations Centers. These will interface with the other operational layers to enable the Network to fulfill its function and satisfy the requirements of the users.

SUMMARY

Design of the Mars Network is proceeding apace, with an ongoing emphasis on maintaining alignment with the rapidly evolving exploration mission set. Despite this atmosphere of change, a number of network functions and elements have been specified, such as communications and navigation provided by satellites in low Mars orbit and in Areostationary orbit. Operations of individual network elements, as well as functions provided by each individual element are reasonably well understood and currently being subjected to a detailed design process. For other aspects of the network, such as network-level functions and operations, design considerations have been identified. Chief among these is the need to provide an efficient, and largely autonomous, switching mechanism to coordinate the demands of the user assets. Other considerations include data management, demand vs. scheduled access, prioritization and qualities of service, handoffs and multiple satellites in view. The network topology must support the switched links among the Mars surface, Mars orbit, Mars approach trajectory, Earth and inter-network cross-links. Further, this topology must evolve to accommodate planned enhancements and unplanned anomalies. A suite of protocols is being developed to guide the implementation of the network functions and topology. These factors will combine to create a virtual presence, first at Mars, and ultimately, via the "Interplanetary Internet," throughout the solar system.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. NASA Strategic Plan (NASA Policy Directive 1000.1): http://www.hq.nasa.gov/office/nsp/

2. Cesarone, R. J., Hastrup, R. C., Bell, D. J., Lyons, D. T. and Nelson, K. G., "Architectural Design for a Mars Communications & Navigation Orbital Infrastructure," AAS/AIAA Astrodynamics Specialist Conference paper AAS 99-300, Girdwood, Alaska, August 16–19, 1999.

3. Bell, D. J., Cesarone, R. J., Ely, T. A., Edwards, C. D. and Townes, S. A., "Mars Network: A Mars Orbiting Communications & Navigation Satellite Constellation," IEEE Aerospace Conference paper 252, Big Sky, Montana, March 18–25, 2000.

4. Hastrup, R. C., Bell, D. J., Cesarone, R. J., Edwards, C. D., Ely, T. A., Guinn, J. R., Rosell, S. N., Srinivasan, J. M. and Townes, S. A., "Mars Network for Enabling Low-Cost Missions," Fourth IAA International Conference on Low-Cost Planetary Missions paper IAA-L-0509, Laurel, Maryland, May 2–5, 2000.

5. Blamont, J., "Using Large Launchers for Small Satellites," JPL Publication 96-26, Introductory Lecture to the 10th Annual AIAA/USU Conference on Small Satellites, Logan, Utah, September 16, 1996.

6. Penzo, P., "Mission Design for Mars Missions Using the Ariane ASAP Launch Capability," AAS/AIAA Space Flight Mechanics Meeting Paper AAS 99-106, Breckenridge, Colorado, February 7–10, 1999.

7. Hastrup, R. C., Cesarone, R. J., Morabito, D. D., and Srinivasan, J. M., "Mars Comm/Nav MicroSat Network," 13th Annual AIAA/USU Conference on Small Satellites Paper SSC 99-VII-5, Logan, Utah, August 23–26, 1999.

8. Greenberg, E., Kazz, G. J., and MacMedan, M. L., "Application of an Extended CCSDS Telecommand Standard for All Mars In-situ Telecommunication links," First ESA Workshop on TT&C, Noordwijk, Netherlands, 24 June 1998.

9. Ely, T. A., Anderson, R. L., Bar-Sever, Y. E., Bell, D. J., Guinn, J. R., Jah, M. K., Kallemeyn, P. H., Levene, E. D., Romans, L. J., and Wu, S-C., "Mars Network Constellation Design Drivers and Strategies," AAS/AIAA Astrodynamics Specialist Conference Paper AAS 99-301, Girdwood, Alaska, August 16–18, 1999.